

Reachability in Augmented Interval Markov Chains

Ventsislav Chonev
IST Austria

Abstract—In this paper we propose augmented interval Markov chains (AIMCs): a generalisation of the familiar interval Markov chains (IMCs) where uncertain transition probabilities are in addition allowed to depend on one another. This new model preserves the flexibility afforded by IMCs for describing stochastic systems where the parameters are unclear, for example due to measurement error, but also allows us to specify transitions with probabilities known to be identical, thereby lending further expressivity.

The focus of this paper is reachability in AIMCs. We study the qualitative, exact quantitative and approximate reachability problem, as well as natural subproblems thereof, and establish several upper and lower bounds for their complexity. We prove the exact reachability problem is at least as hard as the famous square-root sum problem, but, encouragingly, the approximate version lies in NP if the underlying graph is known, whilst the restriction of the exact problem to a constant number of uncertain edges is in P. Finally, we show that uncertainty in the graph structure affects complexity by proving NP-completeness for the qualitative subproblem, in contrast with an easily-obtained upper bound of P for the same subproblem with known graph structure.

I. INTRODUCTION

Discrete-time Markov chains are a well-known stochastic model, one which has been used extensively to reason about software systems [CY95], [HJ94], [RKNP04]. They comprise a finite set of states and a set of transitions labelled with probabilities in such a way that the outgoing transitions from each state form a distribution. They are useful for modelling systems with inherently probabilistic behaviour, as well as for abstracting complexity away from deterministic ones. Thus, it is a long-standing interest of the verification community to develop logics for describing properties concerning reliability of software systems and to devise verification algorithms for these properties on Markov chains and their related generalisations, such as Markov decision processes [Bel57], [Put14].

One well-known such generalisation is motivated by how the assumption of precise knowledge of a Markov chain's transition relation often fails to hold. Indeed, a real-world system's dynamics are rarely known exactly, due to incomplete information or measurement error. The need to model this uncertainty and to reason about robustness under perturbations in stochastic systems naturally gives rise to *interval Markov chains (IMCs)*. In this model, uncertain transition probabilities are constrained to intervals, with two different semantic interpretations. Under the *once-and-for-all* interpretation, the given interval Markov chain is seen as representing an uncountably infinite collection of Markov chains refining it, and

the goal is to determine whether some (or alternatively, all) refinements satisfy a given property. In contrast, the *at-every-step* interpretation exhibits a more game-theoretic flavour by allowing a choice over the outgoing transition probabilities prior to every move. The goal is then to determine strategies which optimise the probability of some property being satisfied. Originally introduced in [JL91], interval Markov chains have recently elicited considerable attention: see for example references [SVA06], [CHS08] and [BLW13], which study the complexity of model checking branching- and linear-time properties, as well as [DLL⁺11], where the focus is on consistency and refinement.

While IMCs are very natural for modelling uncertainty in stochastic dynamics, they lack the expressivity necessary to capture dependencies between transition probabilities arising out of domain-specific knowledge of the underlying real-world system. Such a dependency could state for example that, although the probabilities of some set of transitions are only known to lie within a given interval, they are all identical. Disregarding this information and studying only a dependence-free IMC is impractical, as allowing these transitions to vary independently of one another results in a vastly over-approximated space of possible behaviours.

Therefore, in the present paper we propose *augmented interval Markov chains (AIMCs)*, a generalisation of IMCs which allows for dependencies of this type to be described. We study the effect of this added expressivity through the prism of the (existentially quantified) reachability problem under the once-and-for-all interpretation. Our results are the following. First, we show that the full problem is hard for both the famous square-root sum problem (Theorem 6) and for the class NP (Theorem 3). The former hardness is present even when the underlying graph structure is known and acyclic, whilst the latter arises even in the qualitative subproblem when transition intervals are allowed to include zero, rendering the structure uncertain. Second, assuming known structure, we show the approximate reachability problem to be in NP (Theorem 11). Third, we show that the restriction of the reachability problem to a constant number of uncertain (i.e. interval-valued) transitions is in P (Theorem 4).

II. PRELIMINARIES

A. Markov chains

A *discrete-time Markov chain* or simply *Markov chain (MC)* is a tuple $M = (V, \delta)$ which consists of a finite set of *vertices* or *states* V and a *one-step transition function* $\delta : V^2 \rightarrow [0, 1]$

such that for all $v \in V$, we have $\sum_{u \in V} \delta(v, u) = 1$. For the purposes of specifying Markov chains as inputs to decision problems, we will assume δ is given by a square matrix of rational numbers. The transition function gives rise to a probability measure on V^ω in the usual way. We denote the probability of reaching a vertex t starting from a vertex s in M by $\mathbb{P}^M(s \rightarrow t)$. The *structure* of M is its underlying directed graph, with vertex set V and edge set $E = \{(u, v) \in V^2 : \delta(u, v) \neq 0\}$. Two Markov chains with the same vertex set are said to be *structurally equivalent* if their edge sets are identical.

An *interval Markov chain (IMC)* generalises the notion of a Markov chain. Formally, it is a pair (V, Δ) comprising a vertex set V and a transition function Δ from V^2 to the set $\text{Int}_{[0,1]}$ of intervals contained in $[0, 1]$. For the purposes of representing an input IMC, we will assume that each transition is given by a lower and an upper bound, together with two boolean flags indicating the strictness of the inequalities. A Markov chain $M = (V, \delta)$ is said to *refine* an interval Markov chain $\mathcal{M} = (V, \Delta)$ with the same vertex set if $\delta(u, v) \in \Delta(u, v)$ for all $u, v \in V$. We denote by $[\mathcal{M}]$ the set of Markov chains which refine \mathcal{M} . An IMC's structure is said to be *known* if all elements of $[\mathcal{M}]$ are structurally equivalent. Moreover, if there exists some $\epsilon > 0$ such that for all $M = (V, \delta) \in [\mathcal{M}]$ and all $u, v \in V$, either $\delta(u, v) = 0$ or $\delta(u, v) > \epsilon$, then the IMC's structure is ϵ -*known*. An IMC can have known structure but not ϵ -known structure for example by having an edge labelled with an open interval whose lower bound is 0.

An *augmented interval Markov chain (AIMC)* generalises the notion of an IMC further by equipping it with pairs of edges whose transition probabilities are required to be identical. Formally, an AIMC is a tuple (V, Δ, C) , where (V, Δ) is an IMC and $C \subseteq V^4$ is a set of *edge equality constraints*. A Markov chain (V, δ) is said to refine an AIMC (V, Δ, C) if it refines the IMC (V, Δ) and for each $(u, v, x, y) \in C$, we have $\delta(u, v) = \delta(x, y)$. We extend the notation $[\mathcal{M}]$ to AIMCs for the set of Markov chains refining \mathcal{M} .

The *reachability problem* for AIMCs is the problem of deciding, given an AIMC $\mathcal{M} = (V, \Delta, C)$, an initial vertex $s \in V$, a target vertex $t \in V$, a threshold $\tau \in [0, 1]$ and a relation $\sim \in \{\leq, \geq\}$, whether there exists $M \in [\mathcal{M}]$ such that $\mathbb{P}^M(s \rightarrow t) \sim \tau$. The *qualitative* subproblem is the restriction of the reachability problem to inputs where $\tau \in \{0, 1\}$.

Finally, in the *approximate reachability problem*, we are given a (small) rational number ϵ and a reachability problem instance. If \sim is \geq , our procedure is required to accept if there exists some refining Markov chain with reachability probability greater than $\tau + \epsilon/2$, it is required to reject if all refining Markov chains have reachability probability less than $\tau - \epsilon/2$, and otherwise it is allowed to do anything. Similarly if \sim is \leq . Intuitively, this is a promise problem: in the given instance the optimal reachability probability is guaranteed to be outside the interval $[\tau - \epsilon/2, \tau + \epsilon/2]$.

B. First-order theory of the reals

We denote by \mathcal{L} the first-order language $\mathbb{R}\langle +, \times, 0, 1, <, = \rangle$. Atomic formulas in this language are of the form $P(x_1, \dots, x_n) = 0$ and $P(x_1, \dots, x_n) > 0$ for $P \in \mathbb{Z}[x_1, \dots, x_n]$ a polynomial with integer coefficients. We denote by $\text{Th}(\mathbb{R})$ the *first-order theory of the reals*, that is, the set of all valid sentences in the language \mathcal{L} . Let $\text{Th}^\exists(\mathbb{R})$ be the *existential first-order theory of the reals*, that is, the set of all valid sentences in the existential fragment of \mathcal{L} . A celebrated result [Tar51] is that \mathcal{L} admits quantifier elimination: each formula $\phi_1(\bar{x})$ in \mathcal{L} is equivalent to some effectively computable formula $\phi_2(\bar{x})$ which uses no quantifiers. This immediately entails the decidability of $\text{Th}(\mathbb{R})$. Tarski's original result had non-elementary complexity, but improvements followed, culminating in the detailed analysis of [Ren92]:

- Theorem 1.** 1) $\text{Th}(\mathbb{R})$ is complete for **2-EXPTIME**.
 2) $\text{Th}^\exists(\mathbb{R})$ is decidable in **PSPACE**.
 3) If $m \in \mathbb{N}$ is a fixed constant and we consider only existential sentences where the number of variables is bounded above by m , then validity is decidable in **P**.

We denote by $\exists\mathbb{R}$ the class, introduced in [ŠŠ11], which lies between **NP** and **PSPACE** and comprises all problems reducible in polynomial time to the problem of deciding membership in $\text{Th}^\exists(\mathbb{R})$.

C. Square-root sum problem

The *square-root sum* problem is the decision problem where, given $r_1, \dots, r_m, k \in \mathbb{N}$, one must determine whether $\sqrt{r_1} + \dots + \sqrt{r_m} \geq k$. Originally posed in [O'R81], this problem arises naturally in computational geometry and other contexts involving Euclidean distance. Its exact complexity is open. Membership in **PSPACE** is straightforward via a reduction to the existential theory of the reals. Later this was sharpened in [ABKPM09] to **PosSLP**, the complexity class whose complete problem is deciding whether a division-free arithmetic circuit represents a positive number. This class was introduced and bounded above by the fourth level of the counting hierarchy **CH** in the same paper. However, containment of the square-root sum problem in **NP** is a long-standing open question, originally posed in [GGJ76], and the only obstacle to proving membership in **NP** for the exact Euclidean travelling salesman problem. This highlights a difference between the familiar integer model of computation and the Blum-Shub-Smale Real RAM model [BSS89], under which the square-root sum is decidable in polynomial time [Tiw92]. See also [EY09] for more background.

III. QUALITATIVE CASE

In this section, we will focus on the qualitative reachability problem for AIMCs. We show that, whilst membership in **P** is straightforward when the underlying graph is known, uncertainty in the structure renders the qualitative problem **NP**-complete.

Theorem 2. *The qualitative reachability problem for AIMCs with known structure is in P.*

Proof. Let the given AIMC be \mathcal{M} and s, t the initial and target vertices, respectively. Since the structure $G = (V, E)$ of \mathcal{M} is known, the qualitative reachability problem can be solved simply using standard graph analysis techniques on G . More precisely, for any $M \in [\mathcal{M}]$, $\mathbb{P}^M(s \rightarrow t) = 1$ if and only if there is no path in G which starts in s , does not enter t and ends in a bottom strongly connected component which does not contain t . Similarly, $\mathbb{P}^M(s \rightarrow t) = 0$ if and only if there is no path from s to t in G . \square

Theorem 3. *The qualitative reachability problem for AIMCs is NP-complete.*

Proof. Membership in NP is straightforward. The equivalence classes of $[\mathcal{M}]$ under structure equivalence are at most 2^{n^2} , where n is the number of vertices, since for each pair (u, v) of vertices, either an edge (u, v) is present in the structure or not. This upper bound is exponential in the size of the input. Thus, we can guess the structure of the Markov chain in nondeterministic polynomial time and then proceed to solve an instance of the qualitative reachability problem on an AIMC with known structure in polynomial time by Theorem 2.

We now proceed to show NP-hardness using a reduction from 3-SAT. Suppose we are given a propositional formula φ in 3-CNF:

$$\varphi \equiv \varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_k,$$

where each clause is a disjunction of three literals:

$$\varphi_i \equiv l_{i,1} \vee l_{i,2} \vee l_{i,3}.$$

Let the variables in φ be x_1, \dots, x_m .

Let $\mathcal{M} = (V, \Delta, C)$ be the following AIMC, also depicted in Figure 1. The vertex set has $3m + k + 3$ vertices:

$$\begin{aligned} V = & \{x_1, \dots, x_m, \overline{x}_1, \dots, \overline{x}_m\} \\ & \cup \{\varphi_1, \dots, \varphi_k\} \\ & \cup \{S, F\}, \\ & \cup \{v_0, \dots, v_m\} \end{aligned}$$

that is, one vertex for each possible literal over the given variables, one vertex for each clause, two special sink vertices S, F (success and failure) and $m + 1$ auxiliary vertices. Through a slight abuse of notation, we use x_i, \overline{x}_i to refer both to the literals over the variable x_i and to their corresponding vertices in \mathcal{M} , and similarly, φ_i denotes both the clause in the formula and its corresponding vertex.

The transitions are the following. For all $i \in \{1, \dots, m\}$, we have:

$$\begin{aligned} \Delta(v_{i-1}, x_i) &= \Delta(v_{i-1}, \overline{x}_i) = \Delta(x_i, v_i) = \\ \Delta(x_i, F) &= \Delta(\overline{x}_i, F) = \Delta(\overline{x}_i, v_i) = [0, 1]. \end{aligned}$$

For all $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, 3\}$, we have:

$$\Delta(\varphi_i, l_{i,j}) = [0, 1].$$

For all $i \in \{1, \dots, k\}$,

$$\Delta(v_m, S) = \Delta(v_m, \varphi_i) = \left[\frac{1}{k+1}, \frac{1}{k+1} \right].$$

Finally, $\Delta(S, S) = \Delta(F, F) = [1, 1]$. For all other pairs of vertices u, v , we have $\Delta(u, v) = [0, 0]$.

The edge equality constraints are:

$$C = \bigcup_{i=1, \dots, m} \{ (v_{i-1}, x_i, x_i, v_i), (v_{i-1}, x_i, \overline{x}_i, F) \}$$

Intuitively, the sequence of ‘diamonds’ comprised by v_0, \dots, v_m and the vertices corresponding to literals is a *variable setting gadget*. Choosing transition probabilities $\delta(v_{i-1}, x_i) = \delta(x_i, v_i) = 1$, and hence necessarily $\delta(x_i, F) = 0$, corresponds to setting x_i to true, whereas $\delta(v_{i-1}, \overline{x}_i) = \delta(\overline{x}_i, v_i) = 1$ and $\delta(\overline{x}_i, F) = 0$ corresponds to setting x_i to false. On the other hand, the branching from v_m into $\varphi_1, \dots, \varphi_k$ and the edges from clauses to their literals makes up the *assignment testing gadget*. Assigning non-zero probability to the edge $(\varphi_i, l_{i,j})$ corresponds to selecting the literal $l_{i,j}$ as witness that the clause φ_i is satisfied.

Formally, we claim that there exists a Markov chain $M \in [\mathcal{M}]$ such that $\mathbb{P}^M(v_0 \rightarrow S) = 1$ if and only if φ is satisfiable.

Suppose first that φ is satisfiable and choose some satisfying assignment $\sigma : \{x_1, \dots, x_m\} \rightarrow \{0, 1\}$. Let $M = (V, \delta) \in [\mathcal{M}]$ be the refining Markov chain which assigns the following transition probabilities to the interval-valued edges of \mathcal{M} . First, let

$$\begin{aligned} \delta(v_{i-1}, x_i) &= \delta(x_i, v_i) = \delta(\overline{x}_i, F) = \sigma(x_i), \\ \delta(v_{i-1}, \overline{x}_i) &= \delta(\overline{x}_i, v_i) = \delta(x_i, F) = 1 - \sigma(x_i) \end{aligned}$$

for all $i \in \{1, \dots, m\}$. Second, for each clause φ_i , choose some literal $l_{i,j}$ which is true under σ and set $\delta(\varphi_i, l_{i,j}) = 1$ and consequently $\delta(\varphi_i, l) = 0$ for the other literals l . Now we can observe that the structure of M has two bottom strongly-connected components, namely $\{S\}$ and $\{F\}$, and moreover, F is unreachable from v_0 . Therefore, $\mathbb{P}^M(v_0 \rightarrow S) = 1$.

Conversely, suppose there exists some $M = (V, \delta) \in [\mathcal{M}]$ such that $\mathbb{P}^M(v_0 \rightarrow S) = 1$. We will prove that φ has a satisfying assignment. For each $i \in \{1, \dots, m\}$, write

$$\begin{aligned} p_i &= \delta(v_{i-1}, x_i) = \delta(x_i, v_i) = \delta(\overline{x}_i, F), \\ 1 - p_i &= \delta(v_{i-1}, \overline{x}_i) = \delta(\overline{x}_i, v_i) = \delta(x_i, F). \end{aligned}$$

Notice that

$$\mathbb{P}^M(v_0 x_1 F^\omega) = \mathbb{P}^M(v_0 \overline{x}_1 F^\omega) = p_1(1 - p_1),$$

so we can conclude $p_1 \in \{0, 1\}$, otherwise $\mathbb{P}^M(v_0 \rightarrow S) \neq 1$, a contradiction. If $p_1 = 1$, then

$$\mathbb{P}^M(v_0 x_1 v_1 x_2 F^\omega) = \mathbb{P}^M(v_0 x_1 v_1 \overline{x}_2 F^\omega) = p_2(1 - p_2),$$

whereas if $p_1 = 0$, then

$$\mathbb{P}^M(v_0 \overline{x}_1 v_1 x_2 F^\omega) = \mathbb{P}^M(v_0 \overline{x}_1 v_1 \overline{x}_2 F^\omega) = p_2(1 - p_2).$$

Either way, we must have $p_2 \in \{0, 1\}$ to ensure $\mathbb{P}^M(v_0 \rightarrow S) = 1$. Unrolling this argument further shows $p_i \in \{0, 1\}$ for

all i . In particular, there is exactly one path from v_0 to v_m and it has probability 1. Let σ be the truth assignment $x_i \rightarrow p_i$, we show that σ satisfies φ . Indeed, if some clause φ_i is unsatisfied under σ , then its three literals $l_{i,1}, \dots, l_{i,3}$ are all unsatisfied, so $\delta(l_{i,j}, F) > 0$ for all $j = 1, \dots, 3$. Moreover, for at least one of these three literals, say $l_{i,1}$, we will have $\delta(\varphi_i, l_{i,1}) > 0$, so the path $v_0 \dots v_m \varphi_i l_{i,1} F^\omega$ will have non-zero probability:

$$\mathbb{P}^M(v_0 \dots v_m \varphi_i l_{i,1} F^\omega) = \frac{1}{k+1} \delta(\varphi_i, l_{i,1}) \delta(l_{i,1}, F) \neq 0,$$

which contradicts $\mathbb{P}^M(v_0 \rightarrow S) = 1$. Therefore, σ satisfies φ , which completes the proof of NP-hardness and of the Theorem. \square

IV. CONSTANT NUMBER OF UNCERTAIN EDGES

We now shift our attention to the subproblem of AIMC reachability which arises when the number of interval-valued transitions is fixed, that is, bounded above by some absolute constant. Our result is the following.

Theorem 4. *Fix a constant $N \in \mathbb{N}$. The restriction of the reachability problem for AIMCs to inputs with at most N interval-valued transitions lies in \mathbf{P} . Hence, the approximate reachability problem under the same restriction is also in \mathbf{P} .*

Proof. Let $\mathcal{M} = (V, \Delta, C)$ be the given AIMC and suppose we wish to decide whether there exists $M \in [\mathcal{M}]$ such that $\mathbb{P}^M(s \rightarrow t) \sim \tau$. Let $U \subseteq V$ be the set of vertices which have at least one interval-valued outgoing transition, together with s and t :

$$U = \{s, t\} \cup \{u \in V : \exists v \in V. \Delta(u, v) \text{ is not a singleton}\}.$$

Notice that $|U| \leq N + 2 = \text{const}$. Write $W = V \setminus U$, so that $\{U, W\}$ is a partition of V .

Let \mathbf{x} be a vector of variables, one for each interval-valued transition of \mathcal{M} . For vertices v_1, v_2 , let $\delta(v_1, v_2)$ denote the corresponding variable in \mathbf{x} if the transition (v_1, v_2) is interval-valued, and the only element of the singleton set $\Delta(v_1, v_2)$ otherwise. Let φ_1 be the following propositional formula over the variables \mathbf{x} which captures the set of ‘sensible’ assignments:

$$\begin{aligned} \varphi_1 \equiv & \bigwedge_{v_1 \in V} \sum_{v_2 \in V} \delta(v_1, v_2) = 1 \\ & \wedge \bigwedge_{v_1, v_2 \in V} \delta(v_1, v_2) \in \Delta(v_1, v_2) \cap [0, 1] \\ & \wedge \bigwedge_{(a, b, c, d) \in C} \delta(a, b) = \delta(c, d). \end{aligned}$$

There is clearly a bijection between $[\mathcal{M}]$ and assignments of \mathbf{x} which satisfy φ_1 .

For vertices v_1, v_2 , use the notation $v_1 \rightsquigarrow v_2$ to denote the event ‘ v_2 is reached from v_1 along a path consisting only of vertices in W , with the possible exception of the endpoints v_1, v_2 ’. Notice that for all $u \in U$ and $w \in W$, $\mathbb{P}^M(u \rightsquigarrow$

$w)$ is independent of the choice of $M \in [\mathcal{M}]$. Denote these probabilities by $\alpha(w, u)$. They satisfy the system

$$\bigwedge_{w \in W, u \in U} \alpha(w, u) = \delta(w, u) + \sum_{w' \in W} \delta(w, w') \alpha(w', u),$$

which is linear and therefore easy to solve with Gaussian elimination. Thus, assume that we have computed $\alpha(w, u) \in \mathbb{Q}$ for all $w \in W$ and $u \in U$.

Similarly, for all $u_1, u_2 \in U$, write $\beta(u_1, u_2)$ for the probability of $u_1 \rightsquigarrow u_2$. Notice that $\beta(u_1, u_2)$ is a polynomial of degree at most 1 over the variables \mathbf{x} , given by

$$\beta(u_1, u_2) = \delta(u_1, u_2) + \sum_{w \in W} \delta(u_1, w) \alpha(w, u_2).$$

Thus, assume we have computed symbolically $\beta(u_1, u_2) \in \mathbb{Q}[\mathbf{x}]$ for all $u_1, u_2 \in U$.

Finally, for each $u \in U$, let $y(u)$ be a variable and write \mathbf{y} for the vector of variables $y(u)$ in some order. Consider the following formula in the existential first-order language of the real field:

$$\varphi \equiv \exists \mathbf{x} \exists \mathbf{y}. \varphi_1 \wedge \varphi_2 \wedge \varphi_3,$$

where

$$\begin{aligned} \varphi_2 \equiv & y(t) = 1 \wedge \bigwedge_{u \in U \setminus \{t\}} y(u) = \sum_{u' \in U} \beta(u, u') y(u'), \\ \varphi_3 \equiv & y(s) \sim \tau, \end{aligned}$$

and φ_1 is as above. Intuitively, φ_1 states that the variables \mathbf{x} describe a Markov chain in $[\mathcal{M}]$, φ_2 states that \mathbf{y} gives the reachability probabilities from U to t , and φ_3 states that the reachability probability from s to t meets the required threshold τ . The problem instance is positive if and only if φ is a valid sentence in the existential theory of the reals, which is decidable. Moreover, the formula uses exactly $2|U| \leq 2(N + 2) = \text{const}$ variables, so by Theorem 1, the problem is decidable in polynomial time, as required. \square

Notice that removing the assumption of a constant number of interval-valued transitions only degrades the complexity upper bound, but not the described reduction to the problem of checking membership in $Th^\exists(\mathbb{R})$. As an immediate corollary, we have:

Theorem 5. *The reachability problem and the approximate reachability problem for AIMCs are in $\exists\mathbb{R}$.*

Note that Theorem 5 can be shown much more easily, without the need to consider separately U -vertices and W -vertices as in the proof of Theorem 4. It is sufficient to use one variable per interval-valued transition to capture its transition probability as above and one variable per vertex to express its reachability probability to the target. Then write down an existentially quantified formula with the usual system of equations for reachability in a Markov chain obtained by conditioning on the first step from each vertex. While this easily gives the $\exists\mathbb{R}$ upper bound, it uses at least $|V|$ variables, so it is insufficient for showing membership in \mathbf{P} for the restriction to a constant number of interval-valued transitions.

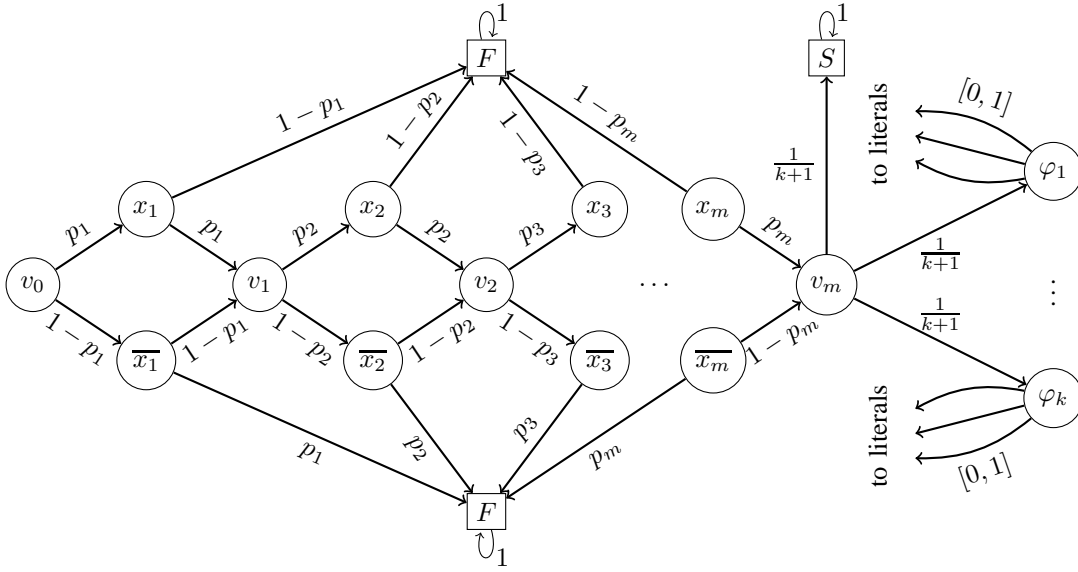


Figure 1. Construction used in Theorem 3 for showing **NP**-hardness of the qualitative AIMC reachability problem. The sink F is duplicated to avoid clutter.

V. HARDNESS FOR SQUARE-ROOT SUM PROBLEM

In this section, we give a lower bound for the AIMC reachability problem. This bound remains in place even when the structure of the AIMC is ϵ -known and acyclic, except for the self-loops on two sink vertices.

Theorem 6. *The AIMC reachability problem is hard for the square-root sum problem, even when the structure of the AIMC is ϵ -known and is acyclic, except for the self-loops on two sink vertices.*

Proof. The reduction is based on the gadget depicted in Figure 2. It is an AIMC with two sinks, S and F (success and failure), each with a self-loop with probability 1, and 12 vertices: $\{a, b_1, \dots, b_4, c_1, \dots, c_4, d_1, d_4, e\}$. The structure is acyclic and comprises four chains leading to S , namely, $ab_1c_1d_1eS$, ab_2c_2S , ab_3c_3S and $ab_4c_4d_4S$. From each vertex other than a and S there is also a transition to F .

The probabilities are as follows. The transition (b_3, c_3) has probability α , whilst (b_1, c_1) , (b_2, c_2) , (b_4, c_4) have probability β , for rationals α, β to be specified later. Consequently, the remaining outgoing transition to F out of each b_i has probability $1 - \alpha$ or $1 - \beta$. The transitions (a, b_i) for $i = 1, \dots, 4$ all have probability $1/4$. Finally, the transitions (c_1, F) , (c_2, F) , (c_3, S) , (c_4, F) , (d_1, e) , (d_4, S) and (e, S) are interval-valued and must all have equal probability in any refining Markov chain. Assign the variable x to the probability of these transitions. The interval to which these transition probabilities are restricted (i.e. the range of x) is to be specified later. Consequently, the remaining transitions $(c_1, d_1), (d_1, F), (e, F), (c_2, S), (c_3, F), (c_4, d_4), (d_4, F)$ are also interval-valued, with probability $1 - x$.

Let M be a positive integer large enough to ensure

$$x^* := \frac{3\sqrt{r}}{2M} \in (0, 1).$$

Then choose a positive integer N large enough, so that

$$\alpha := \frac{4M}{N} \in (0, 1),$$

$$\beta := \frac{16M^3}{27rN} \in (0, 1),$$

$$p_{opt} := \frac{\sqrt{r}}{N} + \frac{\beta}{4} \in (0, 1).$$

Now, a straightforward calculation shows

$$\begin{aligned} \mathbb{P}(a \rightarrow S) &= \mathbb{P}(ab_1c_1d_1eS) + \mathbb{P}(ab_2c_2S) \\ &\quad + \mathbb{P}(ab_3c_3S) + \mathbb{P}(ab_4c_4d_4S) \\ &= \frac{\beta x^2(1-x)}{4} + \frac{\beta(1-x)}{4} \\ &\quad + \frac{\alpha x}{4} + \frac{\beta x(1-x)}{4} \\ &= \frac{\alpha x - \beta x^3 + \beta}{4}. \end{aligned}$$

Analysing the derivative of this cubic, we see that $\mathbb{P}(a \rightarrow S)$ increases on $[0, x^*)$, has its maximum at $x = x^*$ and then decreases on $(x^*, 1]$. This maximum is

$$\frac{\alpha x^* - \beta (x^*)^3 + \beta}{4} = \frac{\sqrt{r}}{N} + \frac{\beta}{4} = p_{opt}.$$

Thus, if we choose some closed interval which contains x^* but not 0 and 1 to be the range of x , then the gadget described thus far will have ϵ -known structure and maximum reachability probability from a to S given by \sqrt{r} scaled by a constant and offset by another constant.

Now, suppose we wish to decide whether $\sqrt{r_1} + \dots + \sqrt{r_m} \geq k$ for given positive integers r_1, \dots, r_m and k . Construct a gadget as above for each r_i . The constants α, N, M are shared across the gadgets, as are the sinks S, F , but each gadget has its own constant β_i in place of β , and its own copy of each non-sink vertex. The edge equality constraints are the same as above within each gadget, and there are no equality constraints across gadgets. Assign a variable x_i to those edges in the i -th gadget which in the description above were labelled x , and choose a range for x_i as described above for x . Finally, add a new initial vertex v_0 , with m equiprobable outgoing transitions to the a -vertices of the gadgets.

In this AIMC, the probability of $v_0 \rightarrow S$ is given by the multivariate polynomial

$$\frac{1}{m} \sum_{i=1}^m \frac{\alpha x_i - \beta_i x_i^3 + \beta_i}{4},$$

whose maximum value on $[0, 1]^m$ is

$$\frac{1}{m} \sum_{i=1}^m \left(\frac{\sqrt{r_i}}{N} + \frac{\beta_i}{4} \right).$$

Therefore, $\sqrt{r_1} + \dots + \sqrt{r_m} \geq k$ if and only if there exists a refining Markov chain of this AIMC with

$$\mathbb{P}(v_0 \rightarrow S) \geq \frac{k}{mN} + \frac{1}{m} \sum_{i=1}^m \frac{\beta_i}{4},$$

so the reduction is complete. \square

Remark 7. It is easy to see that if we are given an acyclic AIMC with the interval-valued edges labelled with variables, the reachability probabilities from all vertices to a single target vertex are multivariate polynomials and can be computed symbolically with a backwards breadth-first search from the target. Then optimising reachability probabilities reduces to optimising the value of a polynomial over given ranges for its variables.

It is interesting to observe that a reduction holds in the other direction as well. Suppose we wish to decide whether there exist values of $x_1 \in I_1, \dots, x_n \in I_n$ such that $P(x_1, \dots, x_n) \geq \tau$ for a given multivariate polynomial P , intervals $I_1, \dots, I_n \subseteq [0, 1]$ and $\tau \in \mathbb{Q}$. Notice that P can easily be written in the form $P(x_1, \dots, x_n) = \beta + N \sum_{i=1}^m \alpha_i Q_i(x_1, \dots, x_n)$, where $N > 0$, $\alpha_1, \dots, \alpha_m \in (0, 1)$ are constants such that $\sum_{i=1}^m \alpha_i \leq 1$, each Q_i is a non-empty product of terms drawn from $\bigcup_{j=1}^n \{x_j, (1 - x_j)\}$, and β is a (possibly negative) constant term. For example, the monomial $-2x_1x_2x_3$ has a negative coefficient, so rewrite it as $2(1 - x_1)x_2x_3 + 2(1 - x_2)x_3 + 2(1 - x_3) - 2$. Do this to all monomials with a negative coefficient, then choose an appropriately large N to obtain the desired form.

Now it is easy to construct an AIMC with two sinks S, F and a designated initial vertex v_0 where the probability of $v_0 \rightarrow S$ is $\sum_{i=1}^m \alpha_i Q_i$. We use a chain to represent each Q_i , and then branch from v_0 into the first vertices of the chains with distribution given by the α_i . There exist values of the

x_i in their appropriate intervals such that $P(x_1, \dots, x_n) \geq \tau$ if and only if there exists a refining Markov chain such that $\mathbb{P}(v_0 \rightarrow S) \geq (\tau - \beta)/N$.

VI. APPROXIMATE CASE

In this section, we focus on the approximate reachability problem for AIMCs. To obtain our upper bound, we will use a result from [Cha12].

Definition 8. If $M_1 = (V, \delta_1)$ and $M_2 = (V, \delta_2)$ are Markov chains with the same vertex set, then their absolute distance is

$$\text{dist}_A(M_1, M_2) = \max_{u, v \in V} \{|\delta_1(u, v) - \delta_2(u, v)|\}.$$

Lemma 9. (Appears in [Cha12].) Let $M_1 = (V, \delta_1)$ and $M_2 = (V, \delta_2)$ be structurally equivalent Markov chains, where $n = |V|$ and for all $u, v \in V$, we have either $\delta_1(u, v) = 0$ or $\delta_1(u, v) \geq \epsilon$. Let also $d \leq \text{dist}_A(M_1, M_2)$ and fix two vertices $s, t \in V$. Then

$$|\mathbb{P}^{M_1}(s \rightarrow t) - \mathbb{P}^{M_2}(s \rightarrow t)| \leq \left(1 + \frac{d}{\epsilon - d}\right)^{2n} - 1.$$

We will also need the following well-known inequality:

Lemma 10. For all $x \geq -1$ and $r \in [0, 1]$, we have

$$(1 + x)^r \leq 1 + rx.$$

Now we proceed to prove our upper bound.

Theorem 11. The approximate reachability problem for AIMCs with ϵ -known structure is in NP.

Proof. Let \mathcal{M} be the given AIMC and let $\epsilon > 0$ be a lower bound on all non-zero transitions across all $M \in [\mathcal{M}]$. Suppose we are solving the maximisation version of the problem: we are given vertices s, t and a rational $\tau > 0$, we must accept if $\mathbb{P}^M(s \rightarrow t) > \tau + \epsilon/2$ for some $M \in [\mathcal{M}]$ and we must reject if $\mathbb{P}^M(s \rightarrow t) < \tau - \epsilon/2$ for all $M \in [\mathcal{M}]$.

Let n be the number of vertices and let

$$d := \epsilon \left(1 - (1 + \epsilon)^{-1/2n}\right).$$

For each interval-valued transition, split its interval into at most $1/d$ intervals of length at most d each. For example, $[l, r]$ partitions into $[l, l + d), [l + d, l + 2d), \dots, [l + kd, r]$, where k is the largest natural number such that $l + kd \leq r$. Call the endpoints defining these subintervals *grid points*. Let $\langle \mathcal{M} \rangle \subseteq [\mathcal{M}]$ be the set of Markov chains refining \mathcal{M} such that the probabilities of all interval-valued transitions are chosen from among the grid points. Observe that for all $M_1 \in [\mathcal{M}]$, there exists $M_2 \in \langle \mathcal{M} \rangle$ such that $\text{dist}_A(M_1, M_2) \leq d$.

Our algorithm showing membership in NP will be the following. We will choose $M \in \langle \mathcal{M} \rangle$ nondeterministically and compute $p := \mathbb{P}^M(s \rightarrow t)$ using Gaussian elimination. Then if $p \geq \tau - \epsilon/2$, we will accept, and otherwise we will reject.

To complete the proof, we need to argue two points. First, that $\langle \mathcal{M} \rangle$ is at most exponentially large in the size of the




Figure 2. Gadget for reduction from square-root sum problem to AIMC reachability.

input, so that M can indeed be guessed in nondeterministic polynomial time. Second, that if for all $M \in \langle \mathcal{M} \rangle$ we have $\mathbb{P}^M(s \rightarrow t) < \tau - \varepsilon/2$, then it is safe to reject, that is, there is no M' with $\mathbb{P}^{M'}(s \rightarrow t) \geq \tau + \varepsilon/2$. (Note that the procedure is obviously correct when it accepts.)

To the first point, we apply Lemma 10 with $x = -\varepsilon/(\varepsilon + 1)$ and $r = 1/2n$:

$$(1 + \varepsilon)^{-1/2n} = \left(1 - \frac{\varepsilon}{1 + \varepsilon}\right)^{1/2n} \leq 1 - \frac{1}{2n} \frac{\varepsilon}{1 + \varepsilon}$$

and hence,

$$d^{-1} = \varepsilon^{-1} \frac{1}{1 - (1 + \varepsilon)^{-1/2n}} \leq \frac{1}{\varepsilon} 2n \frac{1 + \varepsilon}{\varepsilon} = \frac{1}{\varepsilon} 2n \left(1 + \frac{1}{\varepsilon}\right).$$

This upper bound is a polynomial in n , $1/\varepsilon$ and $1/\epsilon$, and hence at most exponential in the length of the input data. Therefore, for each interval-valued transition, we can write down using only polynomially many bits which grid point we wish to use for the probability of that transition. Since the number of transitions is polynomial in the length of the input, it follows that an element of $\langle \mathcal{M} \rangle$ can be specified using only polynomially many bits, as required.

To the second point, consider $M_1, M_2 \in [\mathcal{M}]$ such that $\text{dist}_A(M_1, M_2) \leq d$. Then by Lemma 9, we have

$$\begin{aligned} & |\mathbb{P}^{M_1}(s \rightarrow t) - \mathbb{P}^{M_2}(s \rightarrow t)| \\ & \leq \left(1 + \frac{d}{\epsilon - d}\right)^{2n} - 1 \\ & = \left(\frac{\epsilon}{\epsilon(1 + \varepsilon)^{-1/2n}}\right)^{2n} - 1 \\ & = \varepsilon. \end{aligned}$$

In other words, changing the transition probabilities by at most d does not alter the reachability probability from s to t by more than ε . However, recall that we chose $\langle \mathcal{M} \rangle$ in such a

way that for all $M_1 \in [\mathcal{M}]$, there is some $M_2 \in \langle \mathcal{M} \rangle$ with $\text{dist}_A(M_1, M_2) \leq d$. In particular, if $\mathbb{P}^{M_2}(s \rightarrow t) < \tau - \varepsilon/2$ for all $M_2 \in \langle \mathcal{M} \rangle$, then certainly $\mathbb{P}^{M_1}(s \rightarrow t) < \tau + \varepsilon/2$ for all $M_1 \in [\mathcal{M}]$, so it is safe to reject. This completes the proof. \square

REFERENCES

- [ABKPM09] Eric Allender, Peter Bürgisser, Johan Kjeldgaard-Pedersen, and Peter Bro Miltersen. On the complexity of numerical analysis. *SIAM Journal on Computing*, 38(5):1987–2006, 2009.
- [Bel57] Richard Bellman. A Markovian decision process. Technical report, DTIC Document, 1957.
- [BLW13] Michael Benedikt, Rastislav Lenhardt, and James Worrell. LTL model checking of interval Markov chains. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 32–46. Springer, 2013.
- [BSS89] Lenore Blum, Mike Shub, and Steve Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin (New Series) of the American Mathematical Society*, 21(1):1–46, 1989.
- [Cha12] Krishnendu Chatterjee. Robustness of structurally equivalent concurrent parity games. In *Proceedings of the 15th International Conference on Foundations of Software Science and Computational Structures, FOSSACS'12*, pages 270–285, Berlin, Heidelberg, 2012. Springer-Verlag.
- [CHS08] Krishnendu Chatterjee, Tom Henzinger, and Koushik Sen. Model-checking omega-regular properties of interval Markov chains. In Roberto M. Amadio, editor, *Foundations of Software Science and Computation Structure (FoSSaCS)*, pages 302–317, March 2008.
- [CY95] Costas Courcoubetis and Mihalis Yannakakis. The complexity of probabilistic verification. *Journal of the ACM (JACM)*, 42(4):857–907, 1995.
- [DLL⁺11] Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, and Andrzej Wąsowski. Decision problems for interval Markov chains. In *International Conference on Language and Automata Theory and Applications*, pages 274–285. Springer, 2011.
- [EY09] Kousha Etessami and Mihalis Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *Journal of the ACM (JACM)*, 56(1):1, 2009.

- [GGJ76] M. R. Garey, R. L. Graham, and D. S. Johnson. Some NP-complete geometric problems. In *Proceedings of the Eighth Annual ACM Symposium on Theory of Computing*, STOC '76, pages 10–22, New York, NY, USA, 1976. ACM.
- [HJ94] Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal aspects of computing*, 6(5):512–535, 1994.
- [JL91] Bengt Jonsson and Kim Guldstrand Larsen. Specification and refinement of probabilistic processes. In *Logic in Computer Science, 1991. LICS'91., Proceedings of Sixth Annual IEEE Symposium on*, pages 266–277. IEEE, 1991.
- [O'R81] Joseph O'Rourke. Advanced problem 6369. *Amer. Math. Monthly*, 88(10):769, 1981.
- [Put14] Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [Ren92] James Renegar. On the computational complexity and geometry of the first-order theory of the reals. Part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. *Journal of Symbolic Computation*, 13(3):255 – 299, 1992.
- [RKNP04] Jan J. M. M. Rutten, Marta Kwiatkowska, Gethin Norman, and David Parker. *Mathematical techniques for analyzing concurrent and probabilistic systems*. American Mathematical Soc., 2004.
- [SŠ11] Marcus Schaefer and Daniel Štefankovič. Fixed points, Nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, pages 1–22, 2011.
- [SVA06] Koushik Sen, Mahesh Viswanathan, and Gul Agha. Model-checking Markov chains in the presence of uncertainties. In Holger Hermanns and Jens Palsberg, editors, *TACAS*, pages 394–410, 2006.
- [Tar51] Alfred Tarski. A decision method for elementary algebra and geometry. 1951.
- [Tiw92] Praseon Tiwari. A problem that is easier to solve on the unit-cost algebraic RAM. *Journal of Complexity*, 8(4):393 – 397, 1992.